

AD-A124 941

SUFFICIENT STATISTICS FOR DECENTRALIZED ESTIMATION(U)

1/1

MASSACHUSETTS INST OF TECH CAMBRIDGE LAB FOR
INFORMATION AND DECISION SYSTEMS R R TENNEY NOV 82

UNCLASSIFIED


LIDS-P-1270 N00014-77-C-0532

F/G 9/4

NL

END

FILMED



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

November, 1982

LIDS-P- 1270

(12)

AD A124541

SUFFICIENT STATISTICS FOR DECENTRALIZED ESTIMATION¹

by

Robert R. Tenney²

ABSTRACT

Decentralized estimation problems involve several agents receiving separate noisy observations of a common stochastic process, and each seeks to generate a local estimate of the state of that process. In the general case, these estimates are desired to be consistent in some way, and thus may be jointly penalized with the state via a cost functional to be minimized. In many cases, each agent need only keep track of its local conditional state probability distribution in order to generate the optimal estimates. This paper examines the boundary between problems where this statistic is sufficient and those where it is not; when it is not, the additional information which must be kept appears to have additional structure as illustrated by an example.

DTIC FILE COPY

DTIC
FEB 20 1983
H

¹This work was supported by ONR contract N00014-77-0532C (N041-519).

²M.I.T., L.I.D.S., Rm. 35-213, 77 Massachusetts Ave., Cambridge, MA 02139

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

00 01 02 03 04 05 06 07 08 09 10 11 12

Accension
 M16 M18
 DPC 2:2
 Unannounced
 Justice
 By
 Distribution/
 Availability
 4-11-68
 Page

18

Such problems fall into the class of team theoretic optimization, where the local sensor sites are viewed as separate decision agents acting to achieve some common objective. One of several interesting problems arises when any feedback of the local decisions to the system is ignored-i.e., the problem is one of producing estimates of the system behavior, not controlling it. Applications which exhibit this characteristic include surveillance [2], air traffic control, and multiplatform navigation [3]. The theory which applies to this subclass of problems is that of [4,5], since the lack of feedback and communication (unlike [6,17]) implies a partially nested³ (PN) information structure. The general

³ Each agent, using its own past information, can reconstruct the decisions previously made by any agents which influence its current observation.

approach taken to a PN problem is to reduce it to an equivalent static problem, and this route will be often be followed here.

If both direct feedback and communication are prohibited, the only interesting qualitative issue left is that of second-guessing, where each agent considers the errors others are likely to make (inferred through the relationship of the others' observations to its own information) and adjusts its estimate to be consistent with others. In fact, the need for mutually consistent estimates (decisions) and the resulting information retention requirements of the agents is the major intellectual motivation for this paper.

Thus this work addresses some important applications problems, but also provides a stepping-stone to an understanding of more complex structures. The principal question answered is "when is the local conditional probability distribution enough, when is it not enough, and what more is needed in the latter case?" The contributions are a unified treatment of the decentralized estimation problem, some new (and simpler) proofs and interpretations of existing results, but more importantly an example of what may replace the local state distribution in general dynamical problems.

Subsequent sections specify the problem formulation, establish notation, point out why the decentralized estimation problem becomes trivial if there is not a need for interestimate consistency, and then treat the problem in increasing steps of complexity. First, the static problem is reviewed, then the sequential problem (static system state, but sequential observations which indeed may depend upon an agent's past decisions), and finally the general dynamic case, where the state may evolve randomly in time. It is in the last case where the sufficient statistics start to get interesting, although at least one special case exists.

II. Problem Statement

The specific problem addressed is described here. The general setting is one where the state x of a dynamic system evolves under the influence of a white noise process w . Two agents (generalization to more is straightforward) observe signals y_i which depend only on x , a local, independent white noise process v_i , and a local state x_i . Each generates a decision u_i via a decision rule γ_i which is restricted to be a function of only the past observations and decisions of that agent. These decisions may affect a local dynamic system (local in that its state x_i depends only on itself, a local white noise process w_i , and the local decision u_i), permitting the application of these results to decentralized optimal stopping and search problems (Figure 1). The agents seek to minimize the expected value of a cost function J which is additively separable in time. We seek to find statistics z_1, z_2 and equations determining their behavior such that there exists a pair of decision rules $\hat{\gamma}_1, \hat{\gamma}_2$ with only z_1 (or z_2) as arguments, and which performs as well as the best decision rule which uses all past information. (If the z_i lie in a finite dimensional space, the possible $\hat{\gamma}_i$ may often be characterized by a finite number of parameters, and the original problem reduced to one of parametric optimization.)

The notation is chosen to facilitate the use of various independence assumptions available. Subscripts denote the agent with which a variable is associated. Upper case letters are used to denote sequences, e.g.

$$X_i(s:t) = (x_i(s), \dots, x_i(t)) \quad (2.1)$$

The joint observation and decision are denoted by

$$y(t) = (y_1(t), y_2(t)) \quad u(t) = (u_1(t), u_2(t)) \quad (2.2)$$

The structural assumptions made are stated formally as:




Figure 1. Problem Structure

A1. Open Loop, Markov System:

$$p(x(t+1) | x(t), x_1(t), x_2(t), w(t), w_1(t), w_2(t), U(t), v_1(t), v_2(t), y(t), t) = \quad (2.3)$$

$$p(x(t+1) | x(t), w(t), t)$$

A2. Markov Local Systems:

$p(x_i(t+1) | x_i(t), w_i(t), u_i(t), t)$ completely describes the evolution of $x_i(t+1)$, as in A1.

A3. White Driving Noises:

$w(t)$, $w_1(t)$, and $w_2(t)$ are each independent of all prior random variables.

A4. White Observation Noises:

$v_1(t)$ and $v_2(t)$ are each independent of all prior random variables. Also, $y_i(t)$ is conditionally independent of all prior random variables except⁴ $v_i(t)$, $x(t)$, and $x_i(t)$.

A5. Spatial Independence:

$w(t)$, $w(t)$, $w_2(t)$ are jointly independent; $v_1(t)$ and $v_2(t)$ are jointly independent.

4. $u_i(t-1)$ may be included as part of $x_i(t)$.

A6. Additivity of Objective:

The cost functional J depends only on X , U , X_1 and X_2 , and is additively separable:

$$J(X, U, X_1, X_2) = \sum_{t=1}^T J(x(t), u_1(t), u_2(t), x_1(t), x_2(t), t) \quad (2.4)$$

□

Of these, A1 and A2 simply pose the problem in state space form, and preclude feedback of actions from local systems to the original system, as well as communication between the local systems. A3 and A4 may be relaxed; if colored driving or observation noise is present, state augmentation can be used to reformulate this problem in this framework. A6 is the usual assumption which permits dynamic programming approaches to succeed; if the cost is not additively separable in time, then often the state space can be augmented to make it so (and this is one major motivation for the local dynamic models here, so that the optimal stopping problem can be placed in the present framework.) However, A5 may be of some concern [7], so it is worth pointing out that correlated observation noise can be treated here.

Lemma 1: A problem with

$$p(y_1(t), y_2(t) | x(t)) \neq p(y_1(t) | x(t))p(y_2(t) | x(t)) \quad (2.5)$$

can be reduced to a form satisfying A5.

Proof: Find some statistic $z(t)$ such that

$$p(y_1(t), y_2(t) | x(t), z(t)) = p(y_1(t) | x(t), z(t))p(y_2(t) | x(t), z(t)) \quad (2.6)$$

and augment the state so that $x'(t) = (x(t), z(t))$. Thus (2.6) implies the independence of y_1 and y_2 when conditioned on $x'(t)$. Such a $z(t)$ exists: $z(t) = y(t)$ always works, although statistics of lower dimension may also exist. □

The above formulation is a bit redundant, as the probabilistic representation of state transitions and observation probabilities obviate the need to explicitly consider the w 's and v 's. However, this is the formulation most convenient for the derivations which follow. The redundancy is reduced by assuming that the w 's and v 's are the only primitive sources of randomness, and the above state transition and observation distributions are probabilistic representations of deterministic functions. For example⁵,

$$x(t+1) = f(x(t), w(t), t) \quad \langle m \rangle$$

$$p(x(t+1) | x(t), w(t), t) = \delta(x(t+1); f(x(t), w(t), t))$$

Also, since the general time varying case is being considered, let the first decision be made at $t=1$ so that $w(0)$ can represent initial conditions on the state (and $x(0)$ assumed fixed and known).

In summary, the quantities needed to specify a problem of this type are:

State Dynamics:

$$p(x(t+1) | x(t), w(t), t)$$

$$p(x_i(t+1) | x_i(t), w_i(t), u_i(t), t) \quad i=1,2$$

⁵The δ is either Dirac or Kronecker, depending on the structure of the set in which $x(t+1)$ resides.

Driving noise statistics:

$$\begin{aligned} p(w(0)) & \quad (\text{initial conditions}) \\ p(w(t)) & \\ p(w_i(0)) & \quad (\text{initial conditions}) \quad i=1,2 \\ p(w_i(t)) & \quad i=1,2 \end{aligned}$$

Sensor model:

$$\begin{aligned} p(y_i(t) | x(t), x_i(t), v_i(t), t) & \quad i=1,2 \\ p(v_i(t)) & \quad i=1,2 \end{aligned}$$

Cost:

$$J(x(t), u_1(t), u_2(t), x_1(t), x_2(t), t)$$

The overall objective of the problem is to choose the sequences of decisions $\Gamma_i = \{\gamma_i(.,t), t=1, \dots, T\}$ which are functions⁶ of the local information $I_i(t)$ (note the assumption of perfect local state information)

$$I_i(t) = (y_i(t), u_i(t-1), x_i(t)) \circ I_i(t-1) \quad (2.7)$$

and which minimize

$$J(\Gamma_1, \Gamma_2) = E_{\substack{w \\ w_1 \\ w_2}} \{J(x, u, x_1, x_2)\} \quad (2.8)$$

Since $I_i(t)$ constantly grows in dimension, we seek a smaller but sufficient summary of $I_i(t)$ as a first step in the solution process.

⁶Strictly, these must be measurable functions of $I_i(t)$ so that the expectation in (2.8) is well defined. This and other technical assumptions required for random variables to be well defined will be made implicitly.

It is important that J jointly penalize the decisions in order to require coordination; otherwise the problem becomes much easier.

Lemma 2: If

$$J(x(t), u_1(t), u_1(t), x_1(t), x_2(t), t) = \quad (2.9)$$

$$J_0(x(t), t) + J_1(x(t), u_1(t), x_1(t), t) + J_2(x(t), u_2(t), x_2(t), t)$$

then each agent optimizes J_i separately, independent of the structure of the system pertaining to the other agent. Thus a sufficient statistic for each agent is the local state x_i and the local conditional probability distribution on x , $p(x(t) | y_i(t))$.

Proof: If (2.9) holds, then (2.8) becomes

$$\begin{aligned} E_{\substack{W \\ W_1 \\ W_2}} \sum_{\substack{V_1 \\ V_2}} \{ \sum_{t=1}^T J_0(x(t), t) + \sum_{t=1}^T J_1(x(t), u_1(t), x_1(t), t) \\ + \sum_{t=1}^T \{ J_2(x(t), u_2(t), x_2(t), t) \} \} \end{aligned} \quad (2.10)$$

$$= E_W \{ J_0(x) \} + E_{\substack{W \\ W_1}} \sum_{\substack{V_1 \\ V_2}} \{ J_1(x, u_1, x_1) \} + E_{\substack{W \\ W_2}} \sum_{\substack{V_1 \\ V_2}} \{ J_2(x, u_2, x_2) \} \quad (2.11)$$

by virtue of the independence of U_1 and X_1 from V_2 and W_2 implied by A2-A6 and the structure of Γ_1 . Clearly Γ_1 only affects the second term; hence it is chosen to minimize

$$J_1(\Gamma_1) = E_{\substack{W \\ W_1}} \sum_{\substack{V_1 \\ V_2}} \{ J_1(x, u_1, x_1) \} \quad (2.12)$$

and this is a classical, centralized imperfect state information problem [9]. It is well known that the conditional state distribution is a sufficient statistic for this problem; from the point of view of agent 1, the state of the process external to it which must be considered is $(x(t), x_1(t))$. By assumption, it knows $x_1(t)$ perfectly; thus a sufficient statistic is $x_1(t)$ and the conditional distribution on $x(t)$. A symmetric argument applies to agent 2. \square

Thus we are particularly interested in cases where (2.9) does not hold - where a spatial additive decomposition of the cost does not exist.

Finally, one implication of the above assumptions will be used repeatedly:

Lemma 3: A1-A6, and the restriction on admissible Γ_i , imply that

$$\begin{aligned} p(W_2, V_2, Y_2, U_2, X_2 | W, X, V_1, Y_1, X_1, U_1) \\ = p(W_2, V_2, Y_2, U_2, X_2 | W) \end{aligned} \quad (2.13)$$

Proof: Decompose the first term in (2.13) using Bayes' rule, then invoke A1-A6 and the structure of Γ_i to get

$$p(U_2, X_2 | Y_2, W_2) p(W_2) p(Y_2 | V_2, X) p(V_2) p(X | W) \quad (2.14)$$

and note that W , and only W , appears in the conditioning of (2.14). \square

This summarizes the "spatial Markovness" of the structure embodied by A1-A6, and particularly A5. If one agent knows the entire history of the driving noises for the main system, then it can reconstruct the state sequence (from 2.6), and use this to compute statistics on the random variables of the other agent. No other random variables are associated with

the former can affect this computation, and herein lie the keys to sufficient statistics.

III. Static Problems

The static team estimation problem has been understood for some time; little can be contributed beyond the existing literature [4,5,7,9,10]]. However, as suggested in the introduction, all other problems under consideration can be reduced to this case, so it is worth reviewing to establish the main results.

The static team problem has each agent making one decision based on one observation of the underlying system state. (Figure 2 shows the causality relations). The applicable result is:

Theorem 1: For static teams, the local conditional state distribution is a sufficient statistic for the decision rules.

Proof: Consider the cost

$$J(\Gamma_1, \Gamma_2) = E_{w, v_1, v_2} \{J(x(1), u_1(1), u_2(1))\} \quad (3.1)$$

Fix Γ_2 arbitrarily. If this Γ_2 were optimal, then Γ_1 would minimize

$$E_{w, v_1, v_2} \{J(x(1), u_1(1), u_2(1))\} \quad (3.2)$$

Γ_1 may be defined at each point in its domain separately; here $\gamma(\cdot, 1)$ depends on $I_1(1) = \{y_1(1)\}$. Thus

$$\gamma_1(y_1) = \arg \min_{u_1} E_{w, v_1, v_2} \{J(x(1), u_1, u_2(1)) | y_1\} \quad (3.3)$$

$$= \arg \min_{u_1} E_{w, v_1, v_2} \{ E_{v_1, v_2} \{J(x(1), u_1, u_2(1)) | w, y_1\} | y_1\} \quad (3.4)$$

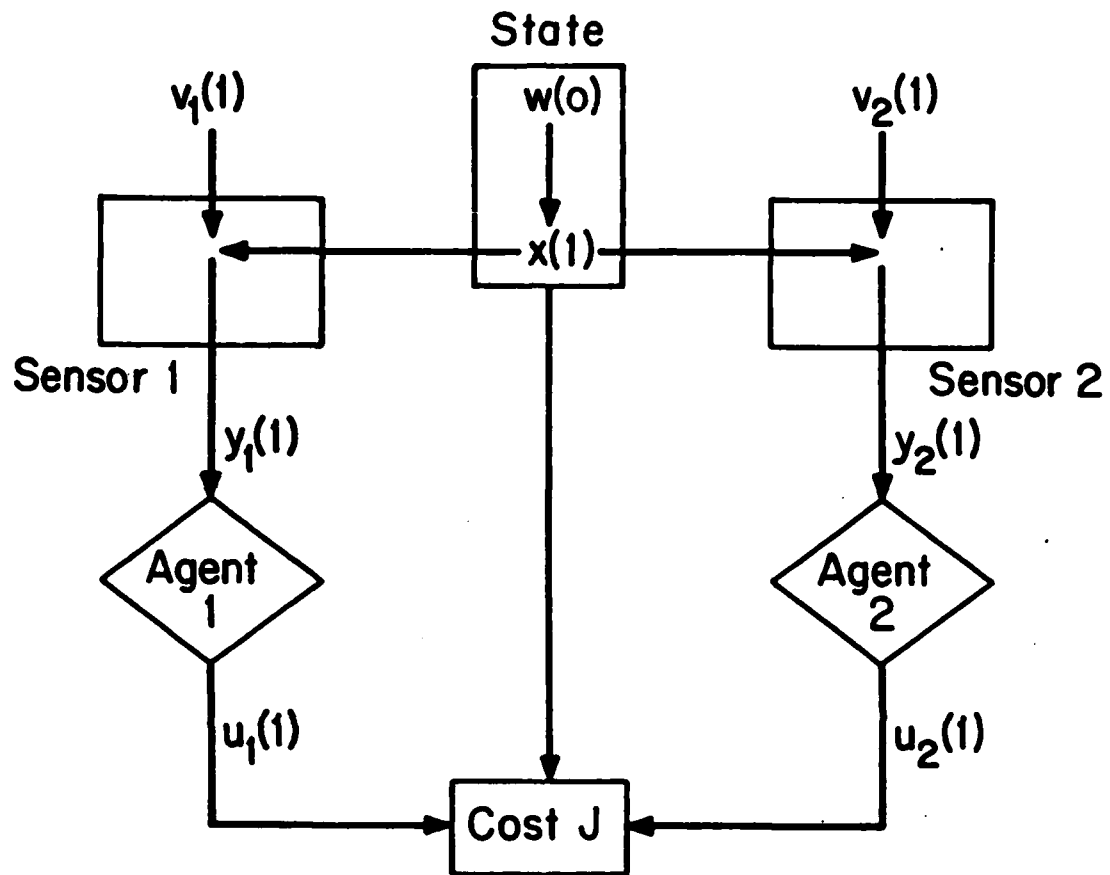


Figure 2. Optimal Solution Structure: Static Case

Since v_1 does not impact J , its expectation may be dropped. The quantity $E\{J(x(1), u_1, u_2(1)) | w, y_1\}$ is independent of y_1 by lemma 3 and the nature of v_2 . Thus it is only a function of $w(0)$ and u_1 , and can be precomputed from Γ_2 ; call it $\hat{J}_1(w(0), u_1)$. Then

$$\gamma_1(y_1) = \arg \min_{u_1} E\{\hat{J}(w, u_1) | y_1\} \quad (3.5)$$

and clearly $p(w(0) | y_1)$ is a sufficient statistic for evaluating this.

□

The above proof exploits the necessary conditions generated by person-by-person-optimality (PBPO) criterion [4], by assuming Γ_2 and deriving properties of γ_1 which must hold for any Γ_2 , including the optimal one. One must be wary of using (3.5) to solve for γ_1 as it is only a necessary condition; here, we have used it only to characterize structural properties of the γ_i .

Example: Suppose $\vec{w}(0) = \vec{x} \in \mathbb{R}^n$ is a vector Gaussian random variable, that $\vec{v}_1 \in \mathbb{R}^{p_1}$ and $\vec{v}_2 \in \mathbb{R}^{p_2}$ are independent Gaussian random variables, and

$$\vec{y}_i = H_i \vec{x} + \vec{v}_i \quad (3.6)$$

are linear observations. Then the solution to this linear, Gaussian (LG) problem is characterized by

Corollary 1a: The conditional mean $E\{x | y_i\}$ is a sufficient statistic for the static LG problem.

Proof: By elementary properties of Gaussian random variables, the sufficient statistic $p(\vec{w}(0) | \vec{y}_i)$ is also Gaussian, and completely defined by its covariance and mean. Its covariance matrix is independent of \vec{y}_i .

Thus the conditional mean⁷ is sufficient for determining $p(\vec{x}|\vec{y}_i)$, and hence u_i . □

Note that this makes no special assumptions on the structure of the cost J . However, when J is jointly quadratic in \vec{x} , \vec{u}_1 , and \vec{u}_2 , γ_i can be found exactly. Let

$$J(\vec{x}, \vec{u}_1, \vec{u}_2) = \frac{1}{2} \begin{bmatrix} \vec{x}^T & \vec{u}_1^T & \vec{u}_2^T \end{bmatrix} \begin{bmatrix} Q_{00} & Q_{01} & Q_{02} \\ Q_{10} & Q_{11} & Q_{12} \\ Q_{20} & Q_{21} & Q_{22} \end{bmatrix} \begin{bmatrix} \vec{x} \\ \vec{u}_1 \\ \vec{u}_2 \end{bmatrix} \quad (3.7)$$

where $\vec{u}_i \in \mathbb{R}^{m_i}$ and the compatibly partitioned matrix Q is symmetric and positive definite. (Note $Q_{21} = Q_{12}^T = 0$ when the cost is spatially separable and Lemma 2 applies.)

Theorem (Radner): The optimal decision rules for the static LQG problem are unique and given by

$$\vec{u}_i = -\underline{G}_i E\{\vec{x}|\vec{y}_i\}$$

where

$$\underline{G}_1 = [Q_{11} - Q_{12}Q_{22}^{-1}Q_{21}]^{-1} [Q_{10} - Q_{12}Q_{22}^{-1}Q_{20}]$$

and symmetrically for \underline{G}_2 .

⁷For reference, $E\{\vec{x}|\vec{y}_i\} = E\{\vec{x}\} + \underline{P} \underline{H}_i^T [\underline{H}_i \underline{P} \underline{H}_i^T + \underline{R}_i]^{-1} (\vec{y}_i - \underline{H}_i E\{\vec{x}\})$, where \underline{P} is the (unconditional) covariance on \vec{x} , \vec{v}_i is zero mean, and \underline{R}_i the covariance matrix \vec{v}_i .

Proof: See [9] or [5]. Note $\underline{Q} > 0$ implies $\underline{Q}_{11} > 0$, and $\underline{Q}_{22} > 0$, and \underline{G}_1 and \underline{G}_2 are well defined.

Thus the static case, as well as the special case of Lemma 2, results in the conditional state distribution being a sufficient statistic.

IV. Sequential Problems

Now we move to a slightly more complex case, where the system state evolves deterministically ($w(1) = w(2) = \dots = w(T-1) = \text{constant}$), but the agents obtain observations and make decisions in a sequence over time. A sufficient statistic must not only supply the requisite information for the current decision, but also must be able to be combined with future observations to generate future sufficient statistics.

First, the problem with a dynamic, but deterministic, evolution of the state $x(t)$ can be reformulated with a time-varying observation structure related to a fixed state - the initial state. If

$$\begin{aligned} x(t+1) &= f(x(t), t) \\ y_i(t) &= h(x(t), v_i(t), t) \end{aligned} \tag{4.1}$$

then defining

$$\begin{aligned} F(x(0), 0) &= x(0) \\ F(x(0), t) &= f(F(x(0), t-1), t) \\ H_i(x(0), t) &= h_i(F(x(0), t), v_i(t), t) \end{aligned} \tag{4.2}$$

is a completely equivalent model relating each $y_i(t)$ to the initial state $x(0) = w(0)$. Note that if a distribution on $w(0)$ is known, an equivalent distribution on $x(t)$ can be found by a straightforward change of variables, but the reverse is true only if $F(\cdot, t)$ is invertible (i.e. one-to-one) (and here lies a clue to the answer of the question posed in the introduction). The remainder of this section will thus focus on $w(0)$ as the only interesting system variable.

That the sequential case is closely related to the static problem can be seen by considering the special case where the local states x_i influence

neither the observations nor the cost.

Corollary 1.b: If $J(x, u_1, u_2, x_1, x_2) = J(x, u_1, u_2)$ and $p(y_i | x, v_i, x_i) = p(y_i | x, v_i)$, then the local conditional initial state distribution $p(w(0) | I_i(t))$ is a sufficient statistic for $\gamma_i(\cdot, t)$.

Proof:

$$\min_{\Gamma_1, \Gamma_2} J(\Gamma_1, \Gamma_2) = \min_{\Gamma_1, \Gamma_2} \sum_{t=1}^T \min_{\gamma_1, \gamma_2} \sum_{w, v_1, v_2} E \{J(w(0), \gamma_1(I_1(t), t), \gamma_2(I_2(t), t), t)\} \quad (4.3)$$

$$= \sum_{t=1}^T \min_{\gamma_1, \gamma_2} \sum_{w, v_1, v_2} E \{J(w(0), \gamma_1(I_1(t), t), \gamma_2(I_2(t), t), t)\} \quad (4.4)$$

because each choice of a decision rule for a particular time t affects exactly one term in the sum. The choices of $\gamma_i(\cdot, t)$ can be separated, and thus the sum and minimization interchanged. From theorem 1, $p(w(0) | I_i(t)) = p(w(0) | y_i(t))$ is a sufficient statistic for γ_i solving the inner (static) team optimization in 4.4. Finally, the sufficient statistic for $\gamma_i(\cdot, t+1)$ can be generated from that for $\gamma_i(\cdot, t)$ and from $y_i(t+1)$ via Bayes' Theorem:

$$p(w(0) | y_i(t+1)) = \frac{p(y_i(t+1) | w(0)) p(w(0) | y_i(t))}{p(y_i(t+1) | y_i(t))} \quad (4.5)$$

where the denominator is directly computable from the terms in the numerator (via summation or integration over $w(0)$).

□

This argument does not readily generalize to the case where local dynamics are present, as the choice of $\gamma_i(\cdot, t)$ influences not only the cost at time t , but also the cost at future times through its effect on x_i (which

may appear directly in the cost, or which influences future observations and hence future costs.) However, the propagation of this effect of each choice of $\gamma_i(\cdot, t)$ is causal, and a nested recursion can be found which, while not a complete separation as in corollary 1.b, provides enough structure to deduce a sufficient statistic.

Theorem 2: For the general sequential problem, where $x(t) = x(0) = w$, a sufficient statistic for each decision rule $\gamma_i(I_i(t), t)$ is the local state $x_i(t)$ combined with the local conditional distribution on w .

Proof: By reverse induction.

Basis: $t=T$. The only term in the cost involving $\gamma_i(I_i(T), T)$ is $J(w, u_1(T), u_2(T), x_1(T), x_2(T))$. Each $\gamma_i(I_i(T), T)$ may be chosen to optimize this term alone. As in Theorem 1, for any Γ_2 , $\gamma_1(I_1(T), T) =$

$$\arg \min_{u_1} E_{w, V_2} \{J(w, u_1, u_2, x_1, x_2) | I_1(T)\} \quad (4.6)$$

$$= \arg \min_{u_1} E_w \{E_{V_2} \{J(w, u_1, u_2, x_1, x_2) | w, I_1(T)\} | I_1(t)\} \quad (4.7)$$

$$= \arg \min_{u_1} E_w \{E_{V_2} (J(w, u_1, u_2, x_1, x_2) | w) | I_1(T)\} \quad (4.8)$$

by Lemma 3. Defining

$$\hat{J}_1(w, u_1, x_1) = E_{V_2} \{J(w, u, u_2, x_1, x_2) | w\} \quad (4.9)$$

it is easy to see that u_i can be chosen to minimize

$$E_w \{ \hat{J}_1(w, u_1, x_1) | I_1(T) \} \quad (4.10)$$

if $p(w | I_1(T))$ and $x_1(T)$ are known. Hence the theorem holds at time T .

Induction: Define $z_i(t) = p(w | I_i(t))$ for convenience. Assume $\gamma_i(z_i(\tau), \tau)$ are fixed for all $\tau = t+1, T$; by the induction hypothesis, such γ_i exist which are equivalent to optimal $\gamma_i(I_i(\tau), \tau)$. Define

$$L(z_1(t+1), z_2(t+1), x_1(t+1), x_2(t+1), w, t+1) =$$

$$E_{w_i(t:T-1)} \sum_{\tau=t+1}^T J(w, x_1, x_2, u_1, u_2, \tau) | I_1(t), I_2(t), w \} \quad (4.11)$$

$$V_i(t+1:T)$$

where the expectation is over the primitive random variables $w_i(\tau)$, $\tau = t, \dots, T-1$, and $v_i(\tau)$, $\tau = t+1, \dots, T$, $i = 1, 2$. Note that this is indeed just a function of z_1, z_2, x_1, x_2 , and w since: the cost at each time is a function of decisions, states, and w ; the states are functions of decisions, prior states, and independent noise; the decisions are functions of the statistics z_i ; the z_i are functions of w and independent noise. Thus all terms in the expectation are, by virtue of A1-A6 and the induction hypothesis, dependent upon $I_1(t), I_2(t)$, and w only through $x_1(t), x_2(t), z_1(t), z_2(t)$, and w - precisely the arguments of L .

Now, consider the choice of $\gamma_1(\cdot, t)$, again with Γ_2 and $\gamma_1(\cdot, \tau)$, $\tau = t+1, \dots, T$, fixed. By the now familiar PBPO arguments, $\gamma_1(I_1(t), t)$ seeks to minimize

$$\begin{aligned}
 & E_{\substack{w \\ w_1 \\ w_2}} \left\{ E_{\substack{v_1 \\ v_2}} \left\{ \sum_{\tau=t}^T J(w, x_1(\tau), x_2(\tau), u_1(\tau), u_2(\tau)) \mid I_1(t) \right\} \right\} \\
 &= E \left\{ E_{\substack{w \\ w_2(t-1)}} \{ J(w, x_1(t), x_2(t), u_1(t), u_2(t)) + \right. \\
 &\quad \left. E_{\substack{v_1 \\ v_2}} \left\{ \sum_{\tau=t+1}^T J(w, x_1(\tau), x_2(\tau), u_1(\tau), u_2(\tau)) \mid I_1(t), w, w_2(t-1), v_2 \right. \right. \\
 &\quad \left. \left. v_1(t+1:T) \right\} \mid w, I_1(t) \} \mid I_1(t) \right\}
 \end{aligned} \tag{4.12}$$

The inner expectation is L , since, $w_2(t-1)$ and $V(t)$ determine $I_2(t)$. By Lemma 3, the middle expectation is independent of $I_1(t)$, since w is included in the conditioning, and we may define

$$\begin{aligned}
 \hat{J}_1(w, x_1(t), z_1(t), u_1) &= \\
 & E_{\substack{w_2(t-1) \\ v_2(t)}} \{ J(w, x_1(t), x_2(t), u_1(t), u_2(t)) + L(z_1(t+1), z_2(t+1), \\
 &\quad x_1(t+1), x_2(t+1), w, t+1) \mid w \}
 \end{aligned} \tag{4.13}$$

The outer expectation and minimization in (4.12) becomes

$$\gamma_1(I_1(t), t) = \arg \min_{u_1} E_{\substack{w}} \{ \hat{J}_1(w, x_1(t), z_1(t), u_1) \mid I_1(t) \} \tag{4.14}$$

for which it is seen that knowledge of $p(w \mid I_1(t))$, $x_1(t)$ and $z_1(t) = p(w \mid I_1(t))$ are sufficient to determine $\gamma_1(\cdot, t)$. □

This result follows directly from the causal structure of the problem. The local state distribution, by Lemma 3, is all that can, and should, be

summarized from $I_1(t)$ to predict the entire behavior of agent 2 both at time t and in the future with all of agent 2's decision rules fixed. This allows agent 1 to predict the impact of 2's decisions on the cost as well as if $I_1(t)$ were all available. $z_1(t) = p(w|I_1(t))$ is all that is necessary to minimize the contribution of $u_1(t)$ to the current cost term, as well as to link $I_1(t)$ to future decisions.

The resulting solution architecture is shown in Figure 3. The local estimators are ordinary Bayesian estimators, each with a structure completely determined by the sensor to which it is attached. Feedback of $x_i(t)$ is required to account for its impact on the observation. The agents now implement $u_i(t) = \gamma_i(z_i(t), x_i(t), t)$ as memoryless decision rules.

The structure of the proof of theorem 2, plus the visualization of Figure 3 which highlights the fact that the statistics $z_i(t)$ evolve as stochastic dynamic systems with inputs w and $x_i(t)$, and driving noise $v_i(t)$, strongly suggests a recursive solution technique, similar to dynamic programming [8], where L plays the role of a cost-to-go function an (z_1, z_2, x_1, x_2, w) that of the state.

This is not quite possible. From figure 3, and the whiteness of $(v_1(t), v_2(t))$, it is clear that the entire system is Markov with a state of (z_1, z_2, x_1, x_2, w) . For a particular choice of $\gamma_1(\cdot, t)$, $\gamma_2(\cdot, t)$, this implies that $p(z_1(t+1), z_2(t+1), x_1(t+1), x_2(t+1), w)$ can be completely determined from $p(z_1(t), z_2(t), x_1(t), x_2(t), w)$ and the γ_i . However, L does not serve to summarize all costs, other than current ones, necessary to choose γ_1 and γ_2 ; the second step of the proof (4.13) required the

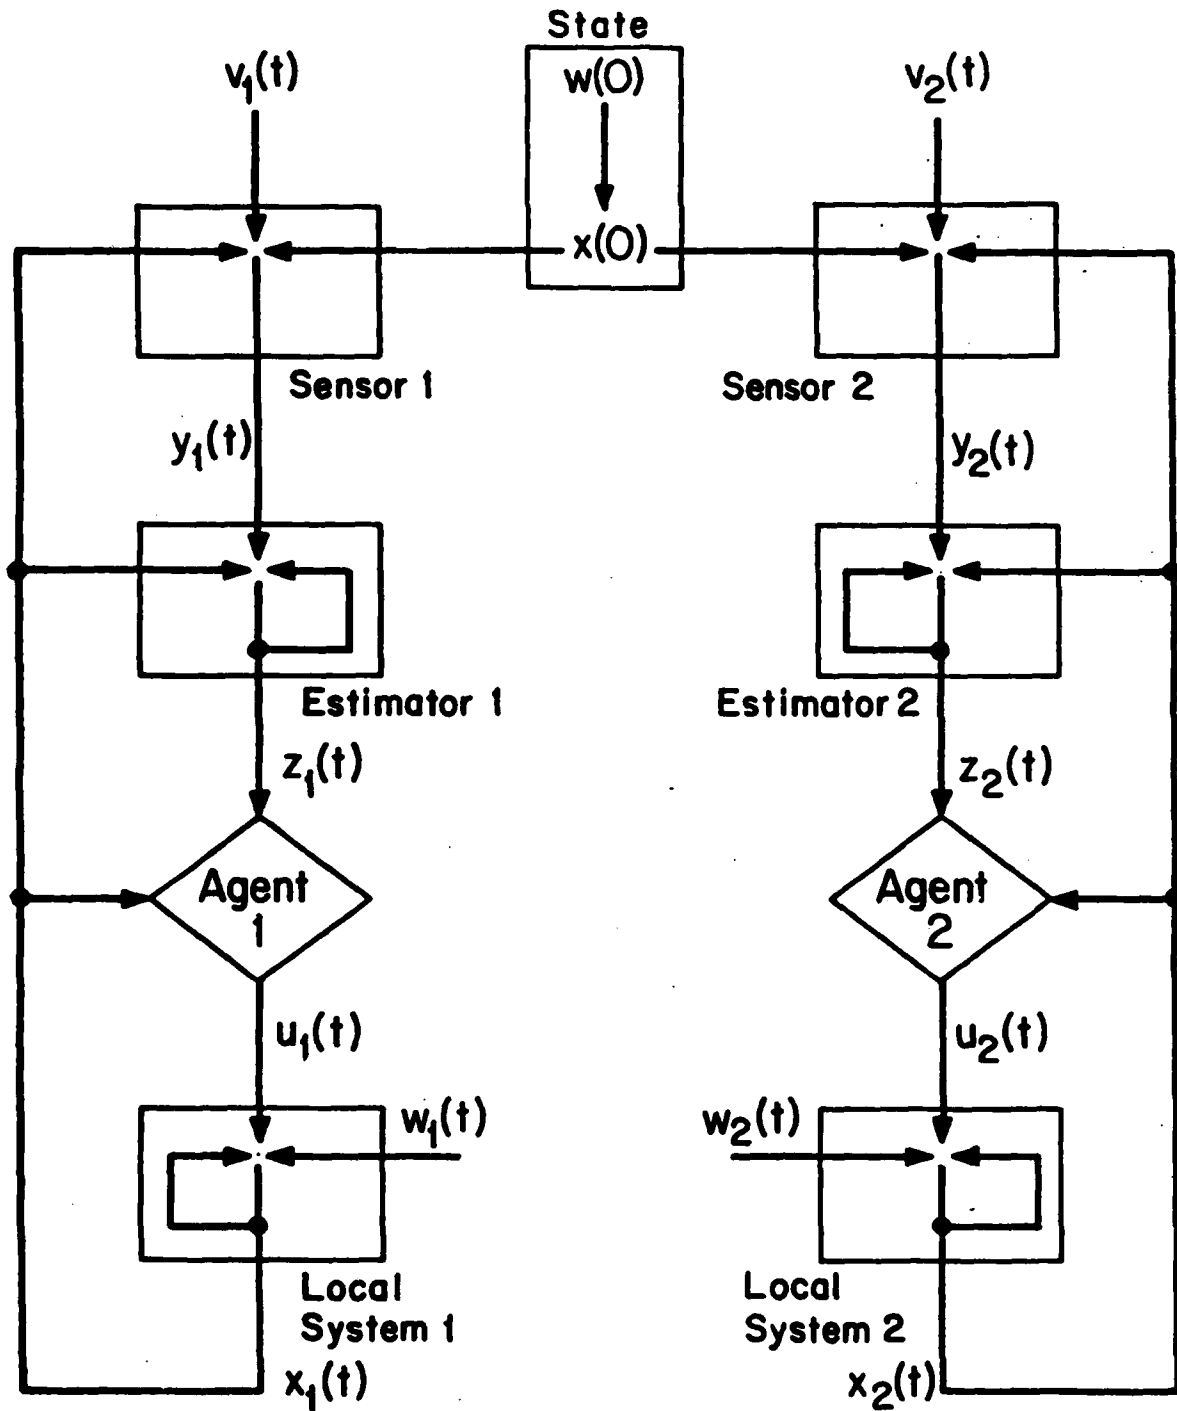


Figure 3. Optimal Solution Structure: Sequential Observations

additional knowledge of $\gamma_2(\cdot, 0), \dots, \gamma_2(\cdot, t)$ in order to exploit PBPO conditions for $\gamma_1(\cdot, t)$. Thus the solution technique resulting from (4.13) would only yield expressions for $\gamma_1(\cdot, t)$ in terms of previous choices of $\gamma_2(\cdot, T)$ - and not separate future from past as in centralized dynamic programming. (The reason is that the choice of γ_1 depends on the $p(u_2(t)|\omega)$, which involves the distribution on $z_2(t)$, which in turn is determined by the prior decision rules of agent 2.)

However, one can get a dynamic programming algorithm by exploiting the joint Markovian structure.

Corollary 2.a: The optimal decision rules for a sequential problem may be determined from a recursion on the joint distribution $p(z_1, z_2, x_1, x_2, w)$:

$$\begin{aligned} \hat{V}(p(z_1, z_2, x_1, x_2, w), T) = \min_{\substack{\gamma_1(\cdot, T), \\ \gamma_2(\cdot, T)}} E \{ J(w, x_1, x_2, u_1, u_2, T) \} \end{aligned} \quad (4.15)$$

and

$$\begin{aligned} \hat{V}(p(z_1, z_2, x_1, x_2, w), t) = \\ \min_{\substack{\gamma_1(\cdot, t) \\ \gamma_2(\cdot, t)}} E \{ J(w, x_1, x_2, z_1, z_2, t) \} \\ + \hat{V}(p(z_1(t+1), z_2(t+1), z_1(t+1), x_1(t+1), x_2(t+1), w), t+1) \end{aligned}$$

where each expectation is over all the random variables inside it, and the probability measure used to evaluate it is that appearing as the argument to \hat{V} . Each $\gamma_i(\cdot, t)$ is restricted to being a function of z_i only.

Proof: The Markovian nature of $p(z_1, z_2, x_1, x_2, w)$ implies that the joint distribution $p(\cdot, \cdot, \cdot, \cdot, \cdot)$ evolves in a purely recursive manner. The deterministic dynamics of this distribution depend only on memoryless control laws of the form specified, independent noise distributions, and local state dynamics; hence it can serve as a dynamic programming state under the conditions specified for the γ_i . \square

This corollary displays the strengths and weaknesses of knowing sufficient statistics z_1 and z_2 . A decentralized decision problem has been reduced to a deterministic dynamic programming problem, from which conclusions as to the behavior of the system under optimal decision policies may be derived. The price paid for this is that of dimensionality - not only are the z_i of higher dimension than the original states, but the dynamic programming is over a probability distribution including the z_i . Thus, while an interesting structurally, this result is unlikely to lead to implementable solution techniques because the double "curse of dimensionality".

Example: Consider the decentralized optimal stopping problem, motivated by [11] and discussed in [12]. The initial state is a binary hypothesis, with known prior distribution $\{p(w=H_A), p(w=H_B)\}$. Each local state x_i is one of three discrete states: continuing (C_i), stopped with H_A declared (A_i), or stopped with H_B declared (B_i). If the local state is

C_i , observations are statistically related to w ; otherwise, they are only noise v_i . Decisions are available which allow the local state C_i to be changed to any local state, but A_i and B_i are trapping states. Initially $x_i(0) = C_i$. Local error penalties are assessed at the terminal time T between the local state and true hypotheses which penalize any event where the local state does not match the true state w . In addition, local data collecting costs are incurred each time the local state is C_i . Finally, to induce coordination, assume that an additional cost is incurred whenever both local states are C_i , thus motivating decision behavior where one agent stops quickly but the other may continue. Application of theorem 2 yields the following characterization of the solution.

Corollary 2.b: A sufficient statistic for the decentralized optimal stopping problem is the local state $x_i \in (A_i, B_i, C_i)$ and the local conditional probability of H_A , $z_i(t) = p(H_A | Y_i(t))$. The optimal decision rule when $x_i = C_i$ is a sequential probability ratio test (SPRT) on $z_i(t)$ with some upper and lower thresholds $\eta_i^1(t)$ and $\eta_i^2(t)$, respectively.

Proof: $z_i(t)$ is sufficient to determine the entire conditional distribution, since w is binary. No effective decision can be made unless $x_i = C_i$. It is straightforward, but tedious, to show that for the cost structure given, any choice of $\gamma_2(\cdot, t)$ leads $J_1(w, x_1 = C_1, z_1(t), u_1)$ to be concave in z_1 when $u_1 = \text{continue}$, and a constant when $u_1 = \text{stop}$ and declare A or B. This implies the SPRT structure. Thus the entire solution is characterized by the $4(T-1)$ parameters $\{\eta_i^0(t), \eta_i^1(t) | i=1,2; t=1,\dots,T-1\}$.

□

Thus the decision rules of the decentralized variation of the optimal stopping problem share the structure of those of the centralized solution, but with different parametric values. Theorem 2 ensures that this is an example of a general phenomenon; since (x_i, z_i) is a sufficient statistic in both the centralized ($i=1$) or decentralized ($i=1,2$) cases, the basic decision structures are identical.

Before concluding this section, the main result of this section can be related to the original question posed in the production by:

Corollary 2c: If the system dynamics are reversible (in 4.1, $f(\cdot, t)$ is one-to-one) in a deterministic, dynamic problem, then $x_i(t)$ and $z_i(t) = p(x(t) | y_i(t))$ is a sufficient statistic for each agent.

Proof: Under these conditions, $p(x(t) | y_i(t))$ completely specifies $p(x(0) | y_i(t))$, which is sufficient by theorem 2. □

V. Dynamic Problems

Consider now the general case of the problem posed in Section II - $x(t)$ evolves as an autonomous Markov process with white driving noise $w(t)$, and each agent receives noisy observations of the state which depend on a local state. This structure is characteristic of many search and surveillance problems, where $x(t)$ models the trajectory of an object, and the two agents are either searching for, or just tracking, the object. The local states model either the trajectory of the search platforms, or the dynamics of the sensor (e.g. pointing a radar).

Following the general procedure of reducing a partially nested team problem to an equivalent static one, some immediate conclusions can be drawn about sufficient statistics in this case.

Theorem 3: Under the basic assumptions A1-A6, a sufficient statistic for each agent in a dynamic estimation problem is the local state $x_i(t)$ in conjunction with the local conditional distribution $p(W(t)|Y_i(t))$ on the driving noise sequence.

Proof: By replacing each w in the proof of theorem 2 with $W(t)$, it is easy to show that $p(W(T-1)|Y_i(t))$ is a sufficient summary of past observations (since $W(T-1)$ can be viewed as an initial, static, state which influences the dynamics in a special way). However, by A3, $p(W(t:T-1)|Y_i(t)) = p(W(t:T-1))$ since $w(t)...w(T-1)$ is white; hence $p(W(T-1)|Y_i(t))$ can be reconstructed from $p(W(t)|Y_i(t))$ and the prior information. □

The result is constructive, but not as helpful computationally as was Theorem 2. Here the sufficient statistic increases in dimension with time - a fact which compounds the dimensionality problem encountered in corollary 2.b. (The sufficient statistic could equally well be taken to be

$p(X(t)|Y_i(t))$ and $x_i(t)$ due to the assumption that w influences y_i and future behavior only through x , and the same problem would exist). However, no claim is made that this is a minimal sufficient statistic; it is possible that other sufficient statistics of fixed dimension can be found.

Example: Suppose the main system is linear

$$\vec{x}(t+1) = \underline{F}(t)\vec{x}(t) + \vec{w}(t) \in \mathbb{R}^n \quad (5.1)$$

with $\vec{w}(t)$ zero-mean and Gaussian. Local observations are linear

$$\vec{y}_i(t) = \underline{H}_i(t)\vec{x}(t) + \vec{v}_i(t) \in \mathbb{R}^{p_i} \quad (5.2)$$

with $\vec{v}_i(t)$ zero-mean and Gaussian. Assume the local states are irrelevant, so each agent seeks to produce directly a local "estimate" $\vec{u}_i(t) \in \mathbb{R}^{m_i}$ to minimize a quadratic cost function as in (3.7). This is the generalization to the dynamic IQG estimation problem of Radner's theorem.

Corollary 3a: For the decentralized IQG estimation problem the local conditional mean on the current state is a sufficient statistic.⁸

Proof. From theorem 3, $p(W(t)|Y_i(t))$ is a sufficient statistic. By elementary properties of Gaussian random variables under linear observations, this distribution is Gaussian specified by a covariance independent of $Y_i(t)$ and conditional mean $E\{\vec{W}(t)|Y_i(t)\}$. By the same argument used in

⁸ Superficially, this seems to contradict the results of [13], where a sufficient statistic was found which increased in dimension with the number of agents. However, that work treated correlated observation noise directly; if Lemma 1 were used to transform that problem to this setting, then it would result here in a new state x of dimension dependent upon the number of agents, and the results are compatible.

corollary 1.b, each $\vec{u}_i(t)$ is chosen to minimize the individual term $E\{J(\vec{x}(t), \vec{u}_1(t), \vec{u}_2(t))\}$. Since J is quadratic, and $\vec{x}(t)$ is a linear combination of the elements of $\vec{W}(t)$, this is a static LQG team problem and Radner's theorem applies (with state $\vec{W}(t)$). In terms of $\vec{u}_1(t)$, $\vec{u}_2(t)$ and $\vec{W}(t)$, this cost, is

$$[\vec{W}^T \vec{u}_1^T \vec{u}_2^T] \begin{bmatrix} \hat{\underline{F}} & 0 & 0 \\ 0 & \underline{I} & 0 \\ 0 & 0 & \underline{I} \end{bmatrix}^T \begin{bmatrix} \underline{Q}_{00} & \underline{Q}_{01} & \underline{Q}_{02} \\ \underline{Q}_{10} & \underline{Q}_{11} & \underline{Q}_{12} \\ \underline{Q}_{20} & \underline{Q}_{21} & \underline{Q}_{22} \end{bmatrix} \begin{bmatrix} \hat{\underline{F}} & 0 & 0 \\ 0 & \underline{I} & 0 \\ 0 & 0 & \underline{I} \end{bmatrix} \begin{bmatrix} \vec{W} \\ \vec{u}_1 \\ \vec{u}_2 \end{bmatrix} \quad (5.3)$$

where

$$\hat{\underline{F}} = [\phi(t,0) : \phi(t,1) : \dots : \phi(t,t-1)] \quad (5.4)$$

and $\phi(t,\tau)$ is the nxn system matrix

$$\underline{\phi}(t,t) = \underline{I} \quad \underline{\phi}(t,\tau) = \underline{\phi}(t,\tau+1)\underline{F}(\tau). \quad (5.5)$$

By Radner's theorem, the optimal decision rule is

$$\vec{u}_1 = -\underline{G}_1(t) E\{\vec{W}(t) | \mathcal{Y}_1(t)\} \quad (5.6)$$

where

$$\begin{aligned} \underline{G}_1(t) &= [\underline{Q}_{11} - \underline{Q}_{12} \underline{Q}_{22}^{-1} \underline{Q}_{21}]^{-1} [\underline{Q}_{10} \hat{\underline{F}} - \underline{Q}_{12} \underline{Q}_{22}^{-1} \underline{Q}_{20} \hat{\underline{F}}] \\ &= \underline{G}_1^* \hat{\underline{F}} \end{aligned} \quad (5.7)$$

The decision rule is then

$$\begin{aligned}\vec{u}_1 &= - \underline{G}_1^* E\{\hat{\underline{F}} \vec{W}(t) | Y_1(t)\} \\ &= - \underline{G}_1^* E\{x(t) | Y_1(t)\}\end{aligned}\tag{5.8}$$

□

This implies that for the dynamic LQG estimation problem, the local Kalman filter estimate is indeed the sufficient statistic. If care is taken to use Lemma 1 to define $\vec{x}(t)$ so that the spatial Markovian property holds, then an elegant result emerges which leads to a computationally feasible solution.

Another interesting point is that the decision rule $\gamma_i(.,t)$, as specified by \underline{G}_1^* , is identical to the rule that would have been used in the static case if $\vec{x}(t)$ were generated alone at time t , with no prior dynamics, and each agent had received an observation $\vec{y}_1(t)$ producing $E\{\vec{x}(t) | \vec{y}(t)\}$ as the conditional mean. Not only does the static nature of the cost yield separation in time of the computation of the decision rules, but the fact that $\vec{x}(t)$ arose as part of a dynamic process does not matter either.

Thus far, several problems have been identified for which the local state and local conditional distribution are sufficient statistics. In the general case, at least so far, only the sufficiency of $p(W(t) | Y_1(t))$ has been shown. Is that as far as we can go, or is the LQG problem indicative of the fact that one more step can be taken to show that $p(x(t) | Y_1(t))$ is sufficient in general?

VI. When the State is Not Enough

Whether or not $p(W(t) | Y_i(t))$ is as far as one can go is best addressed by example. Essentially, $W(t)$ includes information on the entire past trajectory of the process, and we are interested in determining if and when the current state $x(t)$ is enough. Since the system is Markov, and in light of the results thus far, one might conjecture that it is.

Consider a simple, discrete state example. $x(t)$ evolves as a Markov chain, depicted in Figure 4. The states can be interpreted as

- N: normal state
- W: transient warning state
- E: short-lived emergency state

Agent 1 has perfect state information; agent 2 cannot distinguish between N and E, but observes each W (and thus may infer the succeeding E). Each agent makes one of two decisions at each time.

$u_i = 0$ - the system is in N or W

$u_i = 1$ - the system is in the E state.

Penalties are assessed as follows (and added if several apply) :

- (a) 10,000 whenever $u_1(t) \neq u_2(t)$
- (b) 100 whenever $u_i(t) = 1$ and $x(t) \in \{N, W\}$
- (c) 1 whenever $u_i(t) = 0$ and $x(t) = E$.

Thus the agents seek to (a) agree, (b) not generate false alarms, and (c) report emergencies.




Figure 4. Example System Dynamics

A weaker conjecture than the one that $p(x(t)|Y_i(t))$ is sufficient is the following.

Conjecture: If a decision agent has perfect state information in a dynamic, decentralized estimation problem, then its optimal decision rule is a function of the current observation only; i.e., of the current state.

This is certainly true in the single agent case. Consider its consequences in the context of this example.

(1) Cost (a) dominates, as its magnitude relative to the other costs is larger than any ratio of probabilities. Clearly a decision rule exists which never incurs penalty (a), such as $u_1(t) = u_2(t) = 0$ regardless of the data.

(2) Cost (b) is next most significant, and the same decision rule mentioned above also guarantees that (b) will never be incurred. Thus an upper bound on the average cost per stage is $5/19$ - the steady state probability that E is occupied.

(3) By the conjecture and (2), agent 1 must choose $u_1=0$ whenever it sees $x \in \{N,W\}$.


(4) There will be times, long after the most recent W, where agent 2 is not certain whether the state is N or E. By (3) and (1), it must choose $u_2=0$ in these cases.

(5) There is a possibility that the system is in state E in cases such as (4). Agent 2 will be choosing $u_2=0$, so by (1) agent 1 must also choose $u_1=0$.


(6) By the conjecture, since agent 1 must choose $u_1=0$ when $x(t)=E$ in some cases, it must do so in all cases. Thus, if the conjecture holds, the decision rule defined in (2) must be optimal. \square

It is not. By modifying the rule so that $u_1=u_2=1$ every time E is entered immediately after a W, all criteria can be satisfied. Since this is a recurrent event, detectable by both agents, and the penalty (c) is not incurred under the modified rule but is in the original, the modified rule must be strictly better in terms of average cost. However, this is achieved only if agent 1 remembers whether E was entered from W or $\{N,E\}$ - and this is more than just the current state. Thus there are cases where $p(x(t)|Y(t))$ is not enough.

The curious thing about this example is that it is possible to determine exactly what is a sufficient statistic, and that statistic is finite. Consider agent 2; a Bayesian state estimator for it can be in one of three states, $z_2(t)$, representing either E, or W, with probability 1, or the distribution $\{p(N) = .8, p(E) = .2, \text{ and } p(W) = 0\}$. (Note this latter state is trapping until the next W is observed since, for this choice of transition probabilities, the distribution on $\{N,E\}$ achieves steady state after one time step). Agent 1 can infer 2's observations from the original state trajectory, and hence knows its estimator state $z_2(t)$. Viewing the original system and 2's estimator together as a composite, discrete state system, agent 1 sees a system which can be in one of four states (Figure 5.1). Thus agent 1's estimator of the combination tracks both the actual state (upper section of each box), but also the state of agent 2 (lower section).



5.1 System External to Agent 1



5.2 Observable System External to Agent 2

Figure 5. Local Estimation

Similarly, agent 2 can view this extended estimator of agent 1 in combination with the system, and construct a new joint estimator. Surprisingly, it still has three states (Figure 5.2), since states 3 and 4 of agent 1's estimator are not distinguishable to agent 2. Thus finite estimators with states $z_i(t)$ for each agent can be found. When used to augment the system state to $(x(t), z_i(t))$, these produce a composite system the Bayes' estimator of which is the other agent's estimator with state $z_j(t)$. (Moreover, in this case, both z_1 and z_2 are finite.) Note that this is true for any cost function, not just the example cost above; note also that the only change from the computation of $p(x(t) | Y_i(t))$ has been the addition of a state to agent 1's estimator representing the special case where E is entered from W.

The conclusions to be drawn are that examples exist where $p(x(t) | Y_i(t))$ is not a sufficient statistic, but that other sufficient statistics do exist. This example is a bit contrived as the transition probabilities between E and N were chosen so that agent 2's estimator was finite—normally it would be countably infinite. However, there are the suggestions of a procedure for generating sufficient statistics which do apply, but these must wait for a sequel [14].

VII. Conclusions

Theorem 3 is the principal result of this work. In any decentralized problem with the structure specified in section II, each agent must estimate at most the history of system driving noises, which is equivalent to the state trajectory. The intuition behind this is demonstrated by the example in section VI - the past state sequence provides information about the past information received by other agents, and hence allows their decisions to be predicted more accurately than would be possible on the basis of the current state alone.

However, the special cases of section IV, and the LQG dynamic case, show that the local conditional state distributions are sufficient for a number of interesting cases (which include local dynamics), and this reduces the choice of decision rules to seeking memoryless maps from x_i and z_i into u_i . If the infinite time horizon problem were addressed via asymptotic methods, then the search would be further reduced to that of finding a steady-state decision rule of this form (assuming steady-state exists).

The most promising result for future work is the example of Section VI. It illuminates both the nature of the second-guessing phenomenon in decentralized estimation, as well as the fact that the general dynamic case is not always infinitely complex. It is suspected that an algebraic theory of "decentralized realizations" will be required to find structures for the memory of each agent which, taken in conjunction with the system dynamics, produces estimators for another agents which satisfy the symmetric conditions.

References

1. H.L. Van Trees, Detection, Estimation, and Modulation Theory, Part I, Wiley, 1968.
2. M. Athans, W.B. Davenport, Jr., E.R. Ducot, R.R. Tenney, eds. Surveillance and Target Tracking, Proc. 4th MIT/ONR C³ Workshop, MIT/LIDS, Cambridge, Ma.
3. G.F. Gobbini, "Relative Navigation by Means of Passive Ranging", Ph.D. Thesis, MIT Dept. Aero. and Astro. Eng., 1981.
4. Y.C. Ho and K.C. Chu, "Team Decision Theory and Information Structures in Optimal Control Problems, Part I", IEEE Trans. A.C., Vol. AC-17, pp. 15-22, Feb. 1972.
5. K.C. Chu, "Team Decision Theory and Information Structures in Optimal Control Problems, Part II", IEEE Trans. A.C., Vol. AC-17, pp. 22-28, Feb. 1972.
6. H.S. Witsenhausen, "A Counterexample in Stochastic Optimal Control", SIAM J. Contr., Vol. 6, No. 1, 1968.
7. G.S. Lauer and N.R. Sandell, Jr., "Distributed Detection with Waveform Observations: Correlated Observation Processes", Proc. 1982 ACC p. 812, June 1982.
8. D.P. Bertsekas, Dynamic Programming and Stochastic Control, Academic Press, 1976.
9. R. Radner, "Team Decision Problems", Ann. Math. Stat. Vol. 33, No. 3, pp. 857-881, 1962.
10. R.R. Tenney and N.R. Sandell, Jr., "Detection with Distributed Sensors", IEEE Trans. AES, Vol. AES-17, No. 4, July 1981.
11. A. Wald, Sequential Analysis, Wiley, N.Y., 1947.
12. D. Teneketsis, "The Decentralized Quickest Detection Problem", Proc. 21st CDC, Orlando, Fla. Dec. 1982.
13. S.M. Barta, "On Linear Control of Decentralized Stochastic Systems", MIT Ph.D. Thesis, Dept. of EECS, June 1978.
14. R.R. Tenney, "Optimal Decentralized Control of Finite Non-deterministic Processes", to appear in Automatica.
15. B.C. Levy, D.A. Castanon, G.C. Verghase, and A.S. Willsky, "A Scattering Framework for Decentralized Estimation Problems", Proc. 21st DCD, San Diego, CA Dec. 1981.

16. L.K. Ekchian and R.R. Tenney, "Detection Networks", Proc. 21st CDC
Orlando, FLA. Dec. 1982.

END

FILMED

3-83

DTIC